Internet comments as a barometer of public opinion

Elad Oster

The Racah Institute of Physics, The Hebrew University of Jerusalem, Jerusalem 91904, Israel

Erez Gilad

The Unit of Nuclear Engineering, Ben-Gurion University of the Negev, Beer-Sheva 84105, Israel

Alexander Feigel*

Dept. of Applied Physics, Soreq Nuclear Research Center (SNRC), Yavne 81800, Israel

(Dated: March 31, 2015)

Social susceptibility is defined and analyzed using data from CNN news website. The current models of opinion dynamics, voting, and herding in closed communities are extended, and the community's response to the injection of a group with predetermined and permanent opinions is calculated. A method to estimate the values of possible response in Internet communities that follow a specific developing subject is developed. The level of social influence in a community follows from the statistics of responses ("like" and "dislike" votes) to the comments written by the members of the same community. Three real cases of developing news stories are analyzed. We suggest that Internet comments may predict the level of social response similar to a barometer that predicts the intensity of a coming storm in still calm environment.

PACS numbers:

In recent years, governments throughout the Arab world have been overthrown by uprisings that followed the self-immolation of a single person, Mohamed Bouazizi. Similarly, the Occupy Wall Street protest movement was triggered by a single call to action via a social network. Such cases raise an important question: How can an individual possessing no special reputation or authority mobilize an entire community by a single call to stand and fight, while large and professionally organized companies may remain unnoticed? Answering this question will help estimating the appropriate timing and the required size for an initial group to evoke a large-scale social response.

A clear and strong display of personal opinions affects the decision-making processes of others. This phenomenon of social influence may be either positive or negative. Positive social influence facilitates the correlated behavior called herding[1]. Herding contributes significantly to the formation of market prices[2][3], the results of artificial market experiments[4][5][6], traffic flows[7], voting outcomes[8][9], and dynamics of social networks[10][11][12].

Acute herding phenomena, such as social revolutions or financial crises, are extremely difficult to predict, though they are evident when they occur[13]. A parameter, such as temperature in phase transitions, is required to estimate the stability of a community's opinion, i.e. the potential of a small perturbation to culminate in abrupt changes in opinion dynamics. Therefore, to understand the population dynamics prior to a possible transition, it is important to develop a quantitative analysis of herding

as a function of time.

Internet communities are of special interest for the analysis of the herding phenomenon. Individual opinions are widely exposed in binary form of "like" and "dislike" votes ("likes" and "dislikes") over Internet news websites and via social networks. The data span any important event and expose millions of opinions[14]. Simultaneous analysis of a developing news story and the corresponding herding in relevant Internet communities may provide a unique opportunity to study the opinion dynamics in a population as it approaches a critical point and becomes unstable. To the best of our knowledge, the definition and evaluation of the temporal dynamics of herding phenomenon in Internet communities remains a challenge.

In this Article, we estimate the social influence as a function of time in Internet communities that followed any of the following three news stories reported on the CNN website: the Zimmerman trial, Iran Nuclear Negotiations, and the US government shutdown of 2013. We show continuous herding dynamics in all three cases and significant amplification of social influence near the verdict announcement in the Zimmerman case. The method we propose allows for the quantitative estimation of a community response to the injection of a group of non-responsive individuals with predefined opinion. This quantitative analysis is possible due to our novel approach to herding as the conditional probabilities to agree or disagree with other people's opinions. This approach differs from the generally accepted treatment of herding as a topology of social interactions' network[2].

To estimate social susceptibility, we use a specific type of Internet news discussion. Some Internet news websites provide a commentary section where readers can comment and vote (i.e., like or dislike) other readers' comments (see Fig. 1). A reader can usually vote for any

*Electronic address: alexfeigel@gmail.com

Subject	Article$_1$		Article$_k$	
	Like	Dislike	Like	Dislike
Comment 1	$\uparrow_1^1$	$\downarrow_1^1$	$\uparrow_1^k$	$\downarrow_1^k$
Comment 2	$\uparrow_2^1$	$\downarrow_2^1$	$\uparrow_2^k$	$\downarrow_2^k$
$\cdots$	$\cdots$	$\cdots$	$\cdots$	
Comment i	$\uparrow_i^1$	$\downarrow_i^1$	$\uparrow_i^k$	$\downarrow_i^k$

FIG. 1: Internet news and social influence. Consider articles that follow some developing news story. Contrary to printed newspapers, Internet news websites open some articles for commentary by the public and for expressing like or dislike votes for each comment. Comments, together with their likes and dislikes, are written and voted for from both supporters and opponents of the articles statements. Quantitative data of likes ($\uparrow$) and dislikes ($\downarrow$) reveal the conditional probabilities of individual community members to respond positively or negatively to others opinions. These conditional probabilities reflect the level of social influence in the community and allow monitoring the temporal dependence of the level of social influence by following articles on the same subject from different dates.

number of comments, with the restriction of one vote per comment. These data constitute a natural large scale social experiment where the population responds to some external signal (i.e., a comment). A comment, however, is not completely external, but rather created by a community member who responds to the comments of other community members. Consequently, statistics of Internet comments and responses can be used as a measure for mean field opinion dynamics of the corresponding community.

Consider a large population of N individuals who are debating on a subject S and continuously voting in favor of S (up $\uparrow$) or against it (down $\downarrow$). The debate process implies that individuals may change their vote in time. In our model, the interaction between individual i and any other randomly selected individual j is expressed by the fact that the probability per contact of individual i to vote down ($P_\downarrow^{ij}$) depends on the vote of individual j. This conditional probability is given by [15]:

$$P_\downarrow^{ij} = \begin{cases} \alpha_{ij} & \text{if } s_j = 1 \\ \beta_{ij} & \text{if } s_j = 0 \end{cases} = \alpha_{ij}s_j + \beta_{ij}(1-s_j), \quad (1)$$

where s_j is the vote of individual j ($s_j = 1$ for up vote and $s_j = 0$ for down vote) and parameter α_{ij} (β_{ij}) is the probability per contact of individual i voting down given individual j is voting up (down), respectively, regardless of the vote of individual i prior to the interaction with individual j.

In a well mixed homogeneous population, where the number of contacts per individual is $k_i = N$ and $(\alpha_{ij}, \beta_{ij}) = (\alpha, \beta)$, the probability of an individual to
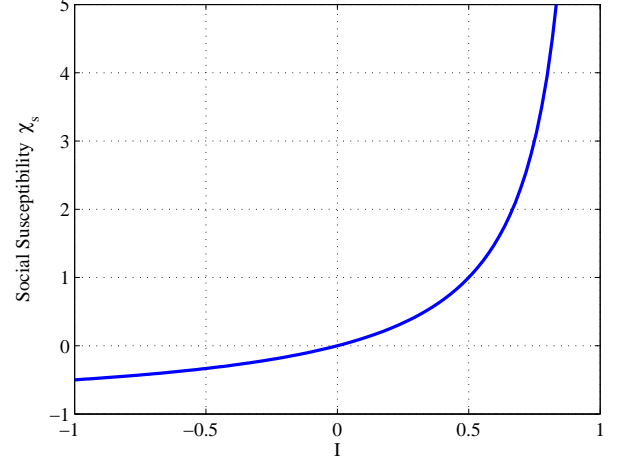

FIG. 2: The social susceptibility χ_s as a function of herding parameter I. Social susceptibility is a measure of how many individuals follow a single one who changes his or her opinion. Thus $\chi_s >> 1$ ($I \approx 1$) makes possible significant social transitions that are initiated by a small group. Social influence vanishes if $\chi_s = 0$. The case of $\chi_s < 0$ corresponds to the populations with negative (antagonistic) social influence.

vote down is

$$P_\downarrow \equiv P_\downarrow^{\,i} = \frac{1}{k_i} \sum_{j=1}^{k_i} P_\downarrow^{\,ij} = \frac{1}{N} \sum_{j=1}^{N} [\alpha s_j + \beta(1-s_j)]. \quad (2)$$

Defining $\gamma \equiv \frac{1}{N}\sum_{i=1}^{N} s_i$ as the mean fraction of individuals who vote up, and noting that mean field assumptions imply $\gamma = P_\uparrow = 1 - P_\downarrow$, Eq. (2) may be written as

$$P_\downarrow = 1 - \gamma = \gamma\alpha + (1-\gamma)\beta, \quad (3)$$

resulting in a steady state expression for γ (the "public opinion") as a function of conditional probabilities

$$\gamma = \frac{1-\beta}{1+\alpha-\beta}. \quad (4)$$

In order to measure social influence, consider a population of N individuals characterized by (α, β), which is perturbed by applying the specific value of mean vote γ_ρ to a fraction $\rho \in [0,1]$ of the population. The new mean vote of the population γ_n is given by

$$\gamma_n = (1-\rho)[\gamma_n(1-\alpha) + (1-\gamma_n)(1-\beta)] + \rho\gamma_\rho. \quad (5)$$

The response function of the population $R(\rho)$ is defined by the fraction of players who flip votes *in response* to the perturbation, i.e. outside the perturbation group. An explicit expression for $R(\rho)$ is obtained using Eq. (5)

$$R(\rho) = \frac{\text{sign}(\gamma_\rho - \gamma)(\gamma_n - \gamma) - \rho|\gamma_\rho - \gamma|}{1-\rho} = \frac{I\rho}{1 - I(1-\rho)}, \quad (6)$$

where $I = \beta - \alpha$. Obviously, the population response function $R(\rho)$ is zero for $\rho = 0$ and for $\alpha = \beta$.

The herding parameter $I \in [-1, 1]$ is a measure of the social influence of one individual on others, because $I = \beta - \alpha$ is the difference of conditional probabilities for correlated and anti-correlated behaviors, see 1. It is similar to herding or percolation parameter $0 \le c \le 1$ from [2]. However, since our definition of the herding parameter I accounts for both positive and negative social influence, it is better suited for analyzing opinion dynamics in binary vote communities.

The social susceptibility χ_s, is defined as

$$\chi_s \equiv \left. \frac{\partial R}{\partial \rho} \right|_{\rho = 0} = \frac{I}{1 - I}, \tag{7}$$

and is the average size of a group whose members follow the change of opinion of a single member (not including the initiating member itself). The size of the perturbation group ρ_{crit} required to convert a population (α, β) to the mean vote of the perturbed group γ_ρ (including the polarized cases $\gamma_\rho = 1, 0$) is obtained by substituting $\gamma_n = \gamma_\rho$ in Eq. (6) and using Eq. (7)

$$\rho_{crit} = \frac{1}{\chi_s} \frac{|\Delta\gamma|}{(1 - |\Delta\gamma|)}, \tag{8}$$

where $\Delta\gamma = \gamma_\rho - \gamma$.

Conditional probabilities (α, β) define herding I, which in turn defines the social stability of the community. To calculate conditional probabilities α and β as a function of likes $\uparrow_i$ and dislikes $\downarrow_i$ votes for comment i of article k (see Fig. 1), we assume that voters and commentators populations are equivalent and that the number of comments and votes is large enough to apply mean field assumption. Consequently, the probabilities for a commentator and a voter to be in favor of the article subject S are both equal to γ. Therefore, the comments should consist of two groups with opposite opinions and relative sizes γ and $1 - \gamma$, respectively.

According to the definition of the conditional probabilities (Eq. 1), the ratio between likes and all responses (likes and dislikes) for a positive comment (to S) is $1 - \alpha$. However, this ratio for a negative comment (to S) equals β since expressing a like vote for a negative comment is equivalent to expressing a dislike vote for the article subject S commented upon. Consequently, the probability of a dislike vote for a comment P_{dislike} is different from the probability $P_\downarrow$ to dislike subject S, as defined in (Eq. 3). Therefore, the probability of a dislike vote for a comment is:

$$P_{\text{dislike}} = \alpha\gamma + (1 - \beta)(1 - \gamma) = 2(1 - \gamma)\gamma. \tag{9}$$

The result is invariant under the transformation $\gamma \to 1 - \gamma$, reflecting the uncertainty regarding the opinion of the Internet article itself. Hence, the division of the comments into two groups with contrasting opinions does not reveal the opinions themselves. Since χ_s is invariant

under the transformation $\gamma \to 1 - \gamma$, we arbitrarily chose $\gamma > 0.5$. An interesting consequence of Eq. (9) is that $P_{\text{dislike}} < 0.5$, i.e. comments cannot include only dislikes because the community cannot dislike its own opinion.

Calculating of α, β and γ of the community proceeds through iterations. First, all comments are sorted by their like vote fraction. Then, at each step n, the comments are divided into two groups with ratio of γ^n and $1 - \gamma^n$ according to their like vote fraction, where group L receives the γ^n comments with the highest like vote fraction and group D receives all other comments. The population characteristic parameters α^n and β^n are then calculated according to:

$$
\begin{aligned}
1 - \alpha^n &= \sum_{i \in L} \frac{\uparrow_i}{\uparrow_i + \downarrow_i}, \\
\beta^n &= \sum_{i \in D} \frac{\uparrow_i}{\uparrow_i + \downarrow_i}.
\end{aligned}
\tag{10}
$$

A new population mean vote γ^{n+1} is calculated using the values of α^n and β^n:

$$\gamma^{n+1} = \frac{1 - \beta^n}{1 + \alpha^n - \beta^n}. \tag{11}$$

The process is repeated until the convergence of α^n, β^n and γ^n.

TABLE I: The results of social susceptibility calculation for 12 CNN articles from different dates that cover three different events. For each article, the values of conditional probabilities (α, β) and social susceptibility χ_s were calculated.

	Article's topic	Publish Date	α	β	γ	χ_s
1	Zimmerman Trial	24/06/13	0.14	0.53	0.77	0.63
2	Zimmerman Trial	05/07/13	0.12	0.38	0.84	0.34
3	Zimmerman Trial	12/07/13	0.08	0.47	0.87	0.64
4	Zimmerman Trial	13/07/13	0.04	0.57	0.91	1.10
5	Zimmerman Trial	17/07/13	0.04	0.67	0.89	1.70
6	Zimmerman Trial	25/07/13	0.06	0.77	0.79	2.39
7	Iran Nuclear Program	25/10/13	0.16	0.58	0.72	0.74
8	Iran Nuclear Program	23/11/13	0.15	0.56	0.75	0.70
9	Iran Nuclear Program	24/11/13	0.15	0.59	0.73	0.78
10	US Govt. Shutdown	01/10/13	0.16	0.51	0.75	0.53
11	US Govt. Shutdown	02/10/13	0.13	0.51	0.79	0.63
12	US Govt. Shutdown	02/10/13	0.09	0.48	0.85	0.62

The formalism of the analysis of the social influence presented above is applied to news articles published on the CNN website that discuss three different topics. The first story includes six articles, published between June 24th and July 25th, 2013, covering the George Zimmerman Trial [16–21]. These articles cover the legal proceeding, the verdict, and the post-verdict jurors opinions about the trial. The second story includes three articles, published between October 25th and November 25th, 2013, covering the negotiations and signing of the Geneva interim agreement on the Iranian nuclear program [22–24]. The third story includes three articles,
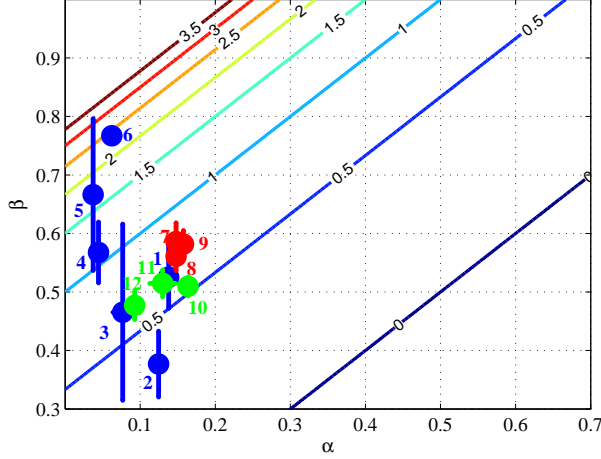
FIG. 3: Social susceptibility χ_s as a function of conditional probabilities (α, β) together with the states of different Internet communities according to the analysis of CNN website articles (numbered dots). The numbers correspond to the cases of Table I that cover the Zimmerman trail, Iran Nuclear Program agreement, and US Government shutdown.
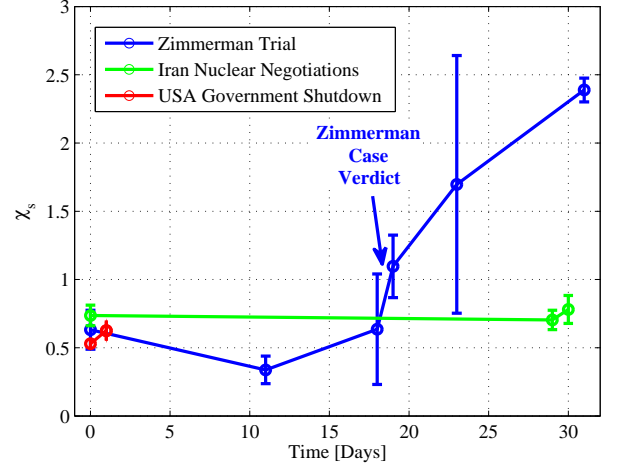

FIG. 4: The social susceptibility χ_s as a function of time. Social susceptibility remains almost constant in the cases of US Government shutdown and Iran Nuclear Program negotiations. This preservation of community state is surprising because there is no reason why the opinion of a population should remain the same for short or for long periods. Even more interesting, however, is that social susceptibility in the case of the Zimmerman trial changes after the verdict is announced on July 13, 2013. Significant social transition becomes possible after the announcement of the verdict.

published on October 1st and 2nd, 2013, covering the US federal government shutdown of that year [25–27]. These articles cover the first day of the shutdown and the White House failing efforts to end it. The results of these analyses are presented in Table I and in Fig. 3.

All three cases exhibit a continuous dynamics in (α, β) space, as shown in Fig. 3. This result is interesting considering that the analysis is applied to different articles, covering different stories, spanning from days to months. It indicates a slow change of opinions in the community.

In the Iran Nuclear Program and US Government shutdown cases, the population's characteristic parameters (α, β) are constant, although they correspond to different CNN articles and, in the case of the Iran Nuclear Program, span one month. This result may also indicate the absence of special events during the observation period.

The social susceptibility level in the Zimmerman trial case changes near the verdict announcement. In the period prior to the verdict day (points 1-3 in Fig. 3), the level of social susceptibility in the population remains almost constant and similar to the social susceptibility in the other cases (i.e., $\chi_s \sim 0.5$), despite the changes in α and β. From the verdict day on (points 4-6), the social susceptibility in the community grows rapidly and the population approaches the singular point $(\alpha, \beta) \to (0, 1)$. It is out of the scope of this work to interpret social phenomena, though the results demonstrate that our method allows to observe the otherwise hidden herding level in a community together with its response to social triggers.

The limitations of our work include the absence of external force, i.e. government control, and lack of interaction topology constrains, such as the prevalence of

near-neighbors interactions. Omitting topological constraints seems to be justified in Internet communities. The same is true regarding forces that shape opinion or add weight to some opinion, such as government control or mass media. We assume that the Internet is still a free zone. The model can be extended to include such force, though there is no clear way to quantify it.

Shortly after the data collection phase for this work was completed, the CNN website has changed its comments policy and the dislike count per comment is no longer displayed. This change made the CNN website articles and comments unsuitable for the above comment analysis procedure, since the main assumption underlying our model, i.e., that both like and dislike vote counts are available to all individuals in the population, is no longer valid. This study demonstrates the potential of both like/dislike votes in estimating the social state of a community and may contribute to the evolving formation of the Internet news format.

To conclude, the developed tools for social influence in Internet communities reveal the previously hidden level of herding and social influence as a function of time in populations. In addition, this work provides a measure for the stability of public opinion in a community and for the size of a group capable to cause critical change in average opinion. The presented method can be compared with other methods and can be extended to other fields such as financial markets[28]. Therefore, this work enables an intriguing comparison of the herding in the same community calculated from different sources, such

as Internet news and financial markets.

[1] A. Banerjee, Quant. J. of Econ. **107**, 797 (1992).
[2] R. Cont and J. Bouchaud, MacroEcon. Dyn. **4**, 170 (2000).
[3] K. Sznajd-Weron and R. Weron, Int. J. of Mod. Phys. C **13**, 115 (2002).
[4] M. J. Salganik, P. S. Dodds, and D. J. Watts, Science **311**, 854 (2006).
[5] C. Borghesi and J.-P. Bouchaud, Quality & Quantity **41**, 557 (2007).
[6] L. Muchnik, S. Aral, and S. J. Taylor, Science **341**, 647 (2013).
[7] D. Helbing, Rev. of Mod. Phys. **73**, 1067 (2001).
[8] R. Durrett, J. P. Gleeson, A. L. Lloyd, P. J. Mucha, F. Shi, D. Sivakoff, J. E. S. Socolar, and C. Varghese, Proc. of Nat. Acad. of Science **109**, 3682 (2012).
[9] J. Fernandez-Gracia, K. Suchecki, J. J. Ramasco, M. San Miguel, and V. M. Eguiluz, Phys. Rev. Lett. **112** (2014).
[10] K. Sznajd-Weron and J. Sznajd, Int. J. of Mod. Phys. C **11**, 1157 (2000).
[11] S. Galam, Phys. A: Stat. Mech. **320**, 571 (2003).
[12] C. Castellano, S. Fortunato, and V. Loreto, Rev. of Mod. Phys. **81**, 591 (2009).
[13] D. Sornette, Phys. Reports **378**, 1 (2003).
[14] R. M. Bond, C. J. Fariss, J. J. Jones, A. D. Kramer, C. Marlow, J. E. Settle, and J. H. Fowler, Nature **489**, 295 (2012).
[15] A. Feigel, J. of Theor. Biol. **254**, 768 (2008).
[16] *Zimmerman opening statements: Expletives and a knock-knock joke*, URL http://edition.cnn.com/2013/06/24/justice/zimmerman-trial/.
[17] *In zimmerman trial, it's a jury of millions*, URL http://edition.cnn.com/2013/07/05/justice/zimmerman-jury-millions/.
[18] *After days of court drama, jurors set to resume deliberating zimmerman's fate*, URL http://edition.cnn.com/2013/07/12/justice/zimmerman-trial/.
[19] *George zimmerman found not guilty of murder in trayvon martin's death*, URL http://edition.cnn.com/2013/07/13/justice/zimmerman-trial/.
[20] *Exclusive: Juror pushes for new laws following zimmerman trial*, URL http://edition.cnn.com/2013/07/17/justice/zimmerman-verdict-aftermath/.
[21] *Zimmerman juror to abc: He 'got away with murder'*, URL http://edition.cnn.com/2013/07/25/justice/zimmerman-juror-b29-interview/.
[22] *Nuclear group: Time iran would need to make uranium for a bomb 'too short'*, URL http://edition.cnn.com/2013/10/25/world/meast/iran-nuclear-report/.
[23] *Obama: Iran nuclear deal limits ability to create nuclear weapons*, URL http://edition.cnn.com/2013/11/23/world/meast/iran-nuclear-talks-geneva/.
[24] *3-decade gridlock broken: The nuclear deal with iran in geneva*, URL http://edition.cnn.com/2013/11/24/world/meast/iran-nuclear-deal/.
[25] *Latest house bid fails as bitter back-and-forth over government shutdown rages*, URL http://edition.cnn.com/2013/10/01/politics/government-shutdown/.
[26] *Progress? obama invites congressional leaders for talks on shutdown.*
[27] *No end in sight to government shutdown after 'unproductive' white house meeting*, URL http://edition.cnn.com/2013/10/02/politics/government-shutdown/.
[28] E. Oster and A. Feigel, submitted (2015).

I. SUPPLEMENTARY MATERIAL

Here we include the detailed procedure to obtain the social influence parameter from Internet discussion data. The algorithms input are two vectors containing the number of likes and dislikes each comments received, $\uparrow_i$ and $\downarrow_i$, respectively (see Fig. s1). The length of these two vectors is the number of comments N, usually few thousands. The output is the resulted population parameters α and β and their error margin.

The initial value of γ_{init} is taken using the probability for a voter to be in favor of a comment - $P_\uparrow^{com}$, which is a measurable parameter given by the ratio of like votes to the total votes:

$$P_\uparrow^{com} = \frac{\sum_i \uparrow_i}{\sum_i \uparrow_i + \sum_i \downarrow_i}. \quad (12)$$

Taking into account (9) and $\alpha = \beta = 1 - \gamma$:

$$P_\uparrow^{com} = \gamma_{init}^2 + (1 - \gamma_{init})^2. \quad (13)$$

The initial value γ_{init} is always chosen to be > 0.5.
Then one proceeds:

1. **Initialization**

   (a) Choose value for the number of voters' threshold: T. Start with $T = 10$.

   (b) From now on, consider only comments above the threshold: $\uparrow_i + \downarrow_i > T$.

   (c) Define the initial value of the mean vote γ by solving the equation:

   $$\frac{\sum_i \uparrow_i}{\sum_i \uparrow_i + \sum_i \downarrow_i} = \gamma^2 + (1 - \gamma)^2.$$

   Take only the solution $\gamma > 0.5$.

2. **Classification of comments**

   (a) Order the comments according to their like vote fraction: $\frac{\uparrow_i}{\uparrow_i + \downarrow_i}$.

   (b) Divide the comments into two groups with ratio of γ and $1 - \gamma$ according to their like vote fraction, i.e., for group L take the γ comments with the highest like vote fraction and for group D take all other comments.

   (c) Calculate the population characteristic parameters α and β:

   $$1 - \alpha = \sum_{i \in L} \frac{\uparrow_i}{\uparrow_i + \downarrow_i},$$

   $$\beta = \sum_{i \in D} \frac{\uparrow_i}{\uparrow_i + \downarrow_i}.$$

   (d) Calculate the new population mean vote γ using the values of α and β:

   $$\gamma = \frac{1 - \beta}{1 + \alpha - \beta}.$$

   (e) Repeat stages (a)-(d) until the values of α and β converge.

3. **Analyzing**

   (a) Increase the threshold for the number of voters T by 1, and repeat stages 1-2.

   (b) End when the number of comments above the threshold N is less than 50.

   (c) The resulted α and β are the weighted mean over all permitted thresholds:

   $$\alpha = \frac{\sum_T \alpha N}{\sum_T N},$$

   $$\beta = \frac{\sum_T \beta N}{\sum_T N}.$$

   (d) The resulting σ_α and σ_β are the equivalent standard deviations over all permitted thresholds.
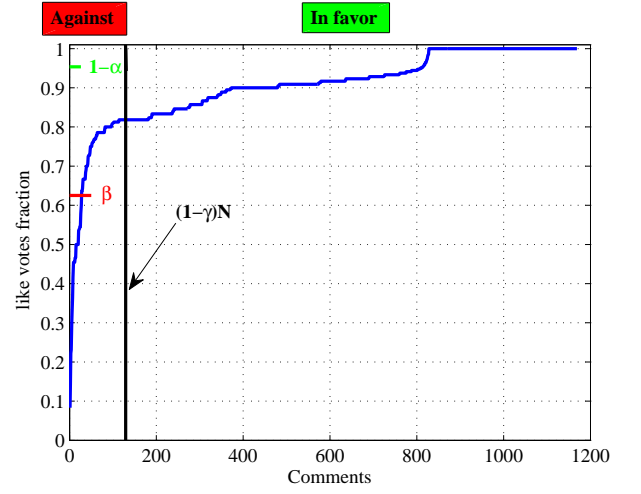

FIG. 5: Like vote fraction distribution for the comments of one CNN article (point 4 in the main article), for threshold value of 10. The black line which is determined by γ, divides the comments into two groups of in favor and against the subject S. The mean like vote fraction of the against and the in favor groups, equal to β and $1 - \alpha$ respectively.

Fig. 5 presents the like vote fraction distribution for the comments of the CNN article announcing the not guilty verdict in the Zimmerman trail (point 4 in the article), for T = 10. The concept of the classification of comments procedure and the way the population parameters α and β are extracted can be well understood in this presentation. For sensitivity of the model to the value of the threshold T see Fig. 6.
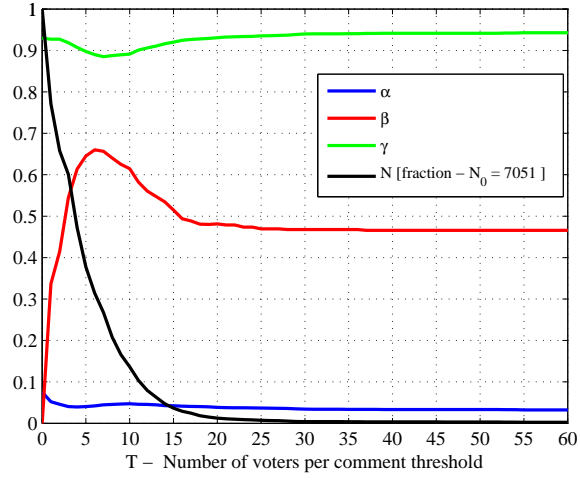
FIG. 6: The sensitivity of the resulting parameters to the threshold value T. The black line represents the number of comments above the threshold and the blue, red, and green lines represent α, β, and γ, respectively.